

09-19-00

A

Please type a plus sign (+) Inside this box

UTILITY PATENT APPLICATION TRANSMITTAL (Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))		Attorney Docket No. 112025-0196
		First Inventor or Application Identifier Darren Kerr et al.
		Title SEQUENCE CONTROL MECHANISM FOR ENABLING OUT OF ORDER CONTEXT PROCESSING
		Express Mail Label No. EL705755824US

APPLICATION ELEMENTS See MPEP chapter 600 concerning utility application contents		Assistant Commissioner for Patents Box Patent Application Washington, DC 20231
--	--	--

1. *Fee Transmittal Form (e.g., PTO/SB/17)
(Submit an original and a duplicate for fee processing)
2. Specification [Total Pages 24]
(preferred arrangement set forth below)

- Descriptive title of the Invention
- Cross References to Related Applications
- Statement Regarding Fed sponsored R & D
- Reference to Microfiche Appendix
- Background of the Invention
- Brief Summary of the Invention
- Brief Description of the Drawings (if filed)
- Detailed Description
- Claim(s)
- Abstract of the Disclosure

3. Drawing(s) [Total Sheets 4]
4. Oath or Declaration [Total Pages 8]

- a. Newly executed (original copy)
- b. Copy from a prior application (37 C.F.R. § 1.63(d))
(for continuation/divisional with Box 17 completed)
[Note Box 5 below]

DELETION OF INVENTOR(S)

Signed statement attached deleting
inventor(s) named in the prior application,
see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).

5. Incorporation By Reference (useable if Box 4b is checked)
 The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered to be part of the disclosure of the accompanying application and is hereby incorporated by reference therein

6. Microfiche Computer Program (Appendix)
7. Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)
- Computer Readable Copy
 - Paper Copy (Identical to computer copy)
 - Statement verifying identity of above copies

- | | |
|---|---|
| ACCOMPANYING APPLICATION PARTS | |
| 8. <input checked="" type="checkbox"/> Assignment Papers (cover sheet & document(s)) | 37 C.F.R. § 3.73(b) |
| 9. <input type="checkbox"/> Statement (when there is an assignee) | <input checked="" type="checkbox"/> Power of Attorney |
| 10. <input type="checkbox"/> English Translation Document (if applicable) | |
| 11. <input type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 | <input type="checkbox"/> Copies of IDS Citations |
| 12. <input type="checkbox"/> Preliminary Amendment | |
| 13. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503)
(Should be specifically itemized) | |
| 14. <input type="checkbox"/> Statement(s) ((PTO/SB/09-12)) | <input type="checkbox"/> Statement filed in prior application,
Status still proper and desired |
| 15. <input type="checkbox"/> Certified Copy of Priority Document(s)
(if foreign priority is claimed) | |
| 16. <input type="checkbox"/> Other: | |

NOTE FOR ITEMS 1 & 14: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28).

17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment:


Continuation Divisional Continuation-in-part (CIP) of prior application No.: /

Prior application Information: Examiner

Group/Art Unit:

18. CORRESPONDENCE ADDRESS

<input checked="" type="checkbox"/> Customer Number or Bar Code Label	<input type="checkbox"/> (Insert Cus)			or <input type="checkbox"/> Correspondence address below <i>(label here)</i>
Name	Charles J. Barbas PATENT TRADEMARK OFFICE			
Address	Cesari and McKenna 88 Black Falcon Avenue			
City	Boston	State	MA	Zip Code
Country	U. S. A.	Telephone	(617) 951-2500	Fax (617) 951-3927

Name (Print/Type)	Charles J. Barbas	Registration No. (Attorney/Agent)	32,959
Signature			Date September 18, 2000

1C682 U.S.P.T.O.
09/18/00
663715

FEE TRANSMITTAL

Patent fees are subject to annual revision on October 1.

These are the fees effective October 1, 1997.

Small Entity payments must be supported by a small entity statement, otherwise large entity fees must be paid. See Forms PTO/SB/09-12. See 37 C.F.R. §§ 1.27 and 1.28.

TOTAL AMOUNT OF PAYMENT (\$) **848**

Complete If Known	
Application Number	Not yet assigned
Filing Date	September 18, 2000
First Named Inventor	Darren Kerr et al.
Examiner Name	Not yet assigned
Group / Art Unit	Not yet assigned
Attorney Docket No.	112025-0196

METHOD OF PAYMENT (check one)

1. The Commissioner is hereby authorized to charge indicated fees and credit any over payments to:

Deposit
Account
Number 03-1237

Deposit
Account
Name Cesari and McKenna, LLP

Charge Any Additional Fee Charge the Issue Fee Set in
Required Under 37 C.F.R. §§1.16 and 1.17 at the Mailing of
the Notice of Allowance

2. Payment Enclosed:

Check Money Other
Order

FEE CALCULATION

1. BASIC FILING FEE

Large Entity Small Entity

Fee Code	Fee (\$)	Fee Code	Fee (\$)	Fee Description	Fee Paid
101	790	201	395	Utility filing fee.	690
106	330	206	165	Design filing fee	
107	540	207	270	Plant filing fee	
108	790	208	395	Reissue filing fee	
114	150	214	75	Provisional filing fee	
SUBTOTAL (1) (\$)		690			

2. EXTRA CLAIM FEES

Extra Fee from
Claims below Fee Paid

Total Claims 20 - 20** = 0 × 18 = 0

Independent Claims 4 - 3** = 1 × 78 = 78

Multiple Dependent = 0

**or number previously paid, if greater; For Reissues, see below

Large Entity Small Entity


Fee Code	Fee (\$)	Fee Code	Fee (\$)	Fee Description
103	22	203	11	Claims in excess of 20
102	82	202	41	Independent claims in excess of 3
104	270	204	135	Multiple dependent claim, if not paid
109	82	209	41	**Reissue independent claims over original patent
110	22	210	11	**Reissue claims in excess of 20 and over original patent
SUBTOTAL (2) (\$)		78		

3. ADDITIONAL FEES

Large Entity	Small Entity	Fee Description	Fee Paid
Fee Code	Fee (\$)	Fee Code (\$)	
105	130	205	65 Fee Surcharge - late filing fee or oath
127	50	227	25 Surcharge - late provisional filing fee or cover sheet
139	130	139	130 Non-English Specification
147	2,520	147	2,520 For filing a request for reexamination
112	920	112	920* Requesting publication of SIR prior to Examiner action
113	1,840	113	1,840*Requesting publication of SIR after Examiner action
115	110	215	55 Extension for reply within first month
116	400	216	200 Extension for reply within second month
117	950	217	475 Extension for reply within third month
118	1,510	218	755 Extension for reply within fourth month
128	2,060	128	1,030 Extension for reply within fifth month
119	310	219	155 Notice of Appeal
120	310	220	155 Filing a brief in support of an appeal
121	270	221	135 Request for oral hearing
138	1,510	138	1,510 Petition to institute a public use proceeding
140	110	240	55 Petition to revive - unavoidable
141	1,320	241	660 Petition to revive - unintentional
142	1,320	242	660 Utility Issue fee (or reissue)
143	450	243	225 Design Issue fee
144	670	244	335 Plant Issue fee
122	130	122	130 Petitions to the Commissioner
123	50	123	50 Petitions related to provisional applications
126	240	126	240 Submission of Information Disclosure Stmt
581	40	581	40 Recording each patent assignment per property (times number of properties) 80
146	790	246	395 Filing a submission after final rejection (37 CFR 1.129(a))
149	790	249	395 For each additional invention to be examined (37 CFR 1.129(b))
Other (specify)			
Other fee (specify)			
SUBTOTAL (3) (\$)			80

*Reduced by Basic Filing Fee Paid

SUBMITTED BY

Complete (if applicable)			
Typed or Printed Name	Charles J. Barbas	Reg. Number	32,959
Signature		Deposit Account User ID	

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re The Application of:)
Darren Kerr et al.)
)
Serial No.: Not yet assigned) Examiner: Not yet assigned
)
Filed: September 18, 2000) Art Unit: Not yet assigned
)
For: SEQUENCE CONTROL MECHA-)
NISM FOR ENABLING OUT OF)
ORDER CONTEXT PROCESSING)
)
)
Cesari and McKenna, LLP
88 Black Falcon Avenue
Boston, MA 02210
September 18, 2000

EXPRESS-MAIL DEPOSIT

“Express Mail” Mailing-Label Number: EL705755824US

The following papers are being deposited with the United States Postal Service
“Express Mail Post Office to Addressee” service pursuant to 37 C.F.R. §1.10:

X Patent Application (24 pages including 20 claims)
X Declaration and Power of Attorney (2) X Fee Transmittal Letter
X Utility Patent Application Transmittal X Assignment With Recordation Form
Letter Cover Sheet (2)
X Formal Drawing (4 sheets) X Check in the amount of \$848.00

UNITED STATES PATENT APPLICATION

of

Darren Kerr

Jeffrey Scott

John Marshall

Kenneth Potter

and

Scott Nellenbach

for a

**SEQUENCE CONTROL MECHANISM FOR ENABLING OUT OF ORDER
CONTEXT PROCESSING**

CROSS-REFERENCE TO RELATED APPLICATIONS

This invention is related to the following copending U.S. Patent Applications:

U.S. Patent Application Serial No. (112025-0197) titled, *Packet Striping Across a Parallel Header Processor*, filed on even date herewith and assigned to the assignee of
5 the present invention; and

U.S. Patent Application Serial No. 09/106,246 titled, *Synchronization and Control System for an Arrayed Processing Engine*, filed on June 29, 1998 and assigned to the as-
signee of the present invention.

FIELD OF THE INVENTION

10 The present invention generally relates to multiprocessor systems and, more spe-
cifically, to out-of-order data processing by processors of an arrayed multiprocessor sys-
tem.

BACKGROUND OF THE INVENTION

A systolic array provides a common approach for increasing processing capacity
15 of a computer system when a problem can be partitioned into discrete units of works. In
the case of a one dimensional systolic array comprising a single “row” of processing ele-
ments or processors, each processor in the array is responsible for executing a distinct set
of instructions on input data before passing it to a next element of the array. To maxi-
mize throughput, the problem is divided such that each processor requires approximately
20 the same amount time to complete its portion of the work. In this way, new input data
can be “pipelined” into the array at a rate equivalent to the processing time of each proc-
essor, with as many units of input data being processed in parallel as there are processors
in the array. Performance can be improved by adding more elements to the array as long
as the problem can continue to be divided into smaller units of work. Once this dividing
25 limit has been reached, processing capacity may be further increased by configuring mul-

tiple rows in parallel, with new input data allocated to the first processor of a next row of the array in sequence.

A symmetric multiprocessor system configured as a systolic array typically guarantees first in, first out (FIFO) ordering of context data processing. As used herein, context data or “context” is defined as an entire packet or, preferably, a header of a packet.

According to FIFO ordering, the contexts processed by the processors of the rows must complete in the order received by the processors before the rows of the array advance.

Each processor is allocated a predetermined time interval or “phase” within which to complete its processing of a context; when each processor completes its context processing within the phase, this control mechanism is sufficient. However if a processor stalls or otherwise cannot complete its processing within the phase interval, all processors of the array stall in order to maintain FIFO ordering. Here, the FIFO ordering control mechanism penalizes both the processors of the row of the stalled processor and the processors of the remaining rows of the multiprocessor array.

For most applications executed by the array, FIFO ordering is not necessary. However, FIFO ordering may be needed to maintain an order of contexts having a dependency among one another; an example of a mechanism used to identify dependencies is a “flow”. A flow is defined as a sequence of packets having the same layer 3 (e.g., Internet Protocol) source and destination addresses, and the same layer 4 (e.g., Transport Control Protocol) port numbers. In addition, the packets of a flow also typically have the same protocol value. The present invention is generally directed to a mechanism that enables selective FIFO ordering of contexts processed by a symmetric multiprocessor system.

SUMMARY OF THE INVENTION

The present invention comprises a sequence control mechanism that enables out-of-order processing of contexts by processors of a symmetric multiprocessor system having a plurality of processors arrayed as a processing engine. The processors of the engine are preferably arrayed as a plurality of rows or clusters embedded between input and output buffers, wherein each cluster of processors is configured to process contexts in a first

in, first out (FIFO) synchronization order. According to the invention, however, the sequence control mechanism allows out-of-order context processing among the clusters of processors, while selectively enforcing FIFO synchronization ordering among those clusters on an as needed basis, i.e., for certain contexts.

5 In the illustrative embodiment, the control mechanism comprises an input sequence controller coupled to the input buffer and an output sequence controller coupled to the output buffer. Each context contains a queue identifier (ID) that uniquely identifies a flow of the context and a sequence number that denotes an order of the context within the flow. A minimum sequence number is used to enforce ordering within a flow having
10 a common queue ID and, to that end, reflects the lowest sequence number of a context for a flow that is active in the processing engine. Synchronization logic of the controllers maintains the minimum (lowest) sequence number for each active flow.

Broadly stated, out-of-order context processing among the clusters is allowed for contexts having different queue IDs, while FIFO synchronization is enforced among the
15 clusters for contexts having the same queue ID. Ordering of contexts associated with a flow is enforced at the output buffer using the queue ID, sequence number and minimum sequence number information associated with each context. That is, the sequence controllers use the information to maintain FIFO synchronization throughout the processing engine, i.e, from the input buffer to the output buffer, for those contexts belonging to a
20 flow and, thus, having the same queue ID.

Advantageously, the inventive sequence control mechanism reduces undesired processing delays among the processors of the arrayed processing engine. Use of the queue ID, sequence number and minimum sequence number information enables the input sequence controller to issue contexts of the same flow to any cluster, thereby exploiting the parallelism inherent in the arrayed processing engine. In addition, the information is used by the output sequence controller to ensure that transmission of a previous context from (off) the processing engine occurs prior to transmission of a subsequent context associated with that flow.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and further advantages of the invention may be better understood by referring to the following description in conjunction with the accompanying drawings in which like reference numbers indicate identical or functionally similar elements:

5 Fig. 1 is a block diagram of a computer network comprising a collection of interconnected communication media and subnetworks attached to a plurality of stations;

Fig. 2 is a schematic block diagram of intermediate station, such as a network switch, that may be advantageously used with the present invention;

10 Fig. 3 is a schematic block diagram of a programmable arrayed processing engine having a plurality of processors configured as clusters;

Fig. 4 is a schematic block diagram of a context adapted for processing by the programmable arrayed processing engine; and

Fig. 5 is a schematic block diagram of a queue content addressable memory that may be advantageously used with the present invention.

15 DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

Fig. 1 is a block diagram of a computer network 100 comprising a collection of interconnected communication media and subnetworks attached to a plurality of stations. The stations are typically computers comprising endstations 102, 112 and intermediate station 200. The intermediate station 200 may be a router or a network switch, whereas the end stations 102, 112 may include personal computers or workstations. The subnetworks generally comprise local area networks (LANs) 110 and 120, although the invention may work advantageously with other communication media configurations such as point-to-point network links. Communication among the stations of the network is typically effected by exchanging discrete data frames or packets between the communicating stations according to a predefined protocol. For the illustrative embodiment described herein, the predefined protocol is the Internet protocol (IP), although the invention could be implemented with other protocols, such as the Internet Packet Exchange, AppleTalk or DECNet protocols.

Fig. 2 is a schematic block diagram of intermediate station 200 that, in the illustrative embodiment, is preferably a network switch. The switch generally performs layer 2 processing functions, such as “cut-through” operations wherein an entire frame does not have to be stored before transfer to a destination; in addition, switch 200 may implement layer 3 forwarding operations. It should be noted, however, that the intermediate station may also be configured as a router to perform layer 3 route processing. A feature of the architecture described herein is the ability to program the station for execution of layer 2, layer 3 or higher-layer operations. Operation of the switch will be described with respect to IP switching of packets, although the switch may be programmed for other applications, such as data encryption.

The switch 200 comprises a plurality of interconnected components including an arrayed processing engine 300, various memories, queuing logic 210 and network port interface cards 240. Operations of these components are preferably synchronously controlled by a clock module 270 although the arrayed elements of the processing engine may be operatively configured to function asynchronously. In the illustrative embodiment, the clock module 270 generates clock signals at a frequency of, e.g., 200 megahertz (i.e., 5 nanosecond clock cycles) and globally distributed them via clock lines to the components of the switch.

The memories generally comprise random access memory storage locations addressable by the processing engine and logic for storing software programs and data structures accessed by the components. An operating system, portions of which are typically resident in memory and executed by the engine, functionally organizes the switch by, *inter alia*, invoking network operations in support of software processes executing on the switch. It will be apparent to those skilled in the art that other memory means, including various computer readable media, may be used for storing and executing program instructions pertaining to the inventive technique and mechanism described herein.

The arrayed processing engine 300 is coupled to a memory partitioned into a plurality of external memory (Ext Mem) resources 280. A feeder chip, hereinafter referred to as a buffer and queuing unit (BQU) 210, comprises data interface circuitry for interconnecting the processing engine with a plurality of interface cards 240 via a selector cir-

cuit 250. Incoming packets to the switch are received at the interface cards 240 and provided to the BQU 210 via the selector 250. The BQU parses a header from each packet and stores it on a data structure, such as linked list, organized as a queue 235 in a queue memory 230. The remaining packet “payload” is stored in a packet memory 220. For
5 each received packet, the BQU builds a *context* for processing by the engine 300. In the illustrative embodiment, the context may comprise the header of each packet, although it may further comprise the entire packet. The BQU 210 provides each context to the processing engine which, after completing processing, returns the context to the BQU where the processed header is appended to the payload of the packet for delivery to the interface
10 cards.

The interface cards 240 may comprise, e.g., OC12, OC48 and Fast Ethernet (FE) ports, each of which includes conventional interface circuitry that may incorporate the signal, electrical and mechanical characteristics, and interchange circuits, needed to interface with the physical media and protocols running over that media. A typical configuration
15 of the switch may include many input/output channels on these interfaces, each of which is associated with one queue 235 in the queue memory 230. The processing engine 300 generally functions as a switching processor that modifies packets and/or headers as the BQU 210 implements queuing operations.

A routing processor 260 executes conventional routing protocols for communication directly with the processing engine 300. The routing protocols generally comprise topological information exchanges between intermediate stations to determine preferred paths through the network based on, e.g., destination IP addresses. These protocols provide information used by the processor 260 to create and maintain routing tables. The tables are loaded into the external partitioned memories 280 as forwarding information
20 base (FIB) tables used by the processing engine to perform forwarding operations. When processing a header in accordance with IP switching, the engine 300 determines where to send the packet by indexing into the FIB using an IP address of the header. Execution of the forwarding operations results in destination media access control (MAC) addresses of
25 the headers being rewritten by the processing engine to identify output ports for the packets.
30

Fig. 3 is a schematic block diagram of the programmable processing engine 300 which comprises an array of processors embedded between input and output buffers with a plurality of interfaces 310 from the array to partitions of an external memory. The external memory stores non-transient data organized within data structures for use in processing the transient data. The non-transient data typically includes “table” data contained in forwarding and routing tables, statistics, access filters, encryption keys and/or queuing information. The transient data enters and exits the engine via 200 MHz 64-bit input and output data interfaces of the BQU 210. A remote processor interface (not shown) provides information, such as instructions and data, from a remote processor to the processors and buffers over a maintenance bus having multiplexed address/data lines.

The processing engine 300 comprises a plurality of processors 350 arrayed into multiple rows and columns that may be further configured as a systolic array. In the illustrative embodiment, the processors are arrayed as eight (8) rows and two (2) columns in an 8x2 arrayed configuration that is embedded between an input buffer 360 and an output buffer 370. However, it should be noted that other arrangements, such as 4x4 or 8x1 arrayed configurations, may be advantageously used with the present invention. The processors of each row are connected to a context memory 330; collectively, these elements of the row are organized as a cluster 345.

Each processor is a customized, single-threaded microcontroller (TMC) 350 having a dense structure that enables implementation of similar processors on an application specific integrated circuit. The present invention may apply to any number of processors within a column of the arrayed engine and, alternatively, to a single processor with multiple threads of execution. The TMC 350 is preferably a pipelined processor that includes, *inter alia*, a plurality of arithmetic logic units (ALUs) and a register file having a plurality of general purpose registers that store intermediate result information processed by the ALUs.

The processors (TMC 0,1) of each cluster 345 execute operations on “transient” data loaded into the context memory 330 by the input buffer 360, whereas the processors of each column operate in parallel to perform substantially the same operation on the transient data, but with a shifted phase. The context memory 330 stores transient “con-

“text” data (e.g., packet/frame data) flowing through the cluster that is unique to a specific process, along with pointers that reference data structures and tables stored in, e.g., Ext Mem 280 for use by the TMC 350.

Each Ext Mem 280 is coupled to an external memory (XRAM) controller 310 which, in the illustrative embodiment, is preferably embodied as a 200 MHz external memory interface coupled to a column of processors. The controller is configured to enable columned processor access to the non-transient data stored in the external column memory. The shared Ext Mem 280 accessed by the processors may further comprise entries of data structures, such as tables, that are constantly updated and accessed by the processors of each column. An example of such a table structure is the FIB table used by the processing engine to perform forwarding operations.

In the illustrative embodiment, the processors of a cluster inherently implement first in, first out (FIFO) ordering primarily because there is no mechanism for bypassing processors within the cluster. Each processor participates in a “software pipeline” phase and if processing by one processor of a cluster stalls (i.e., is delayed), all processors in that cluster are delayed. This arrangement can create undesired dependencies if all phases do not complete within a maximum interval and the contexts are unrelated. In this case, the depth of each cluster directly affects the magnitude of undesired dependencies.

However, each cluster of processors can execute independently with respect to other clusters. Thus, the number of clusters in the processing engine is significant since clusters (rows) of processors increase the concurrency of processing within the engine. Synchronization among the clusters is needed to maintain FIFO ordering. By modifying this synchronization method, undesired dependencies can be avoided by enforcing FIFO synchronization on an as needed basis. Thus, a solution to the problem of a processor requiring additional time to complete its context processing is to allow out-of-order processing among the clusters of processors. That is, if a processor of a cluster stalls and there is no dependency between the active contexts in the stalled cluster and in other clusters of the processing engine, then the contexts in the other clusters may be transmitted from the processing engine once their processing completes. Such concurrency facilitates

the selective FIFO ordering technique described herein to thereby enable out-of-order context processing.

The present invention comprises a sequence control mechanism that enables out-of-order processing of contexts by processors of a symmetric multiprocessor system arrayed as a processing engine to thereby reduce undesired processing delays among those processors. As described herein, the sequence control mechanism (i.e., sequence controller) generally allows out-of-order context processing among the clusters/rows of processors, yet selectively enforces FIFO synchronization ordering among those clusters on an as needed basis, i.e., for certain contexts. To that end, the sequence controller utilizes a queue identifier (ID) and a sequence number assigned to each context to maintain FIFO synchronization throughout the processing engine, i.e., from input buffer 360 to output buffer 370, for those contexts belonging to a same flow. A flow is defined by parameters such as the IP source and destination address, the TCP port numbers and the protocol associated with a packet or group of packets.

In the illustrative embodiment, the BQU 210 assigns a queue ID and, if necessary, a sequence number to each context provided to the input buffer 360. The BQU 210 includes a first circuit 212 configured to execute a conventional hash function that generates a queue ID that uniquely identifies a flow of a context. That is, the flow parameters of a context are transformed by the hash function 212 to a 32-bit queue ID. It should be noted that the hash function 212 may utilize other information that enables identification of dependencies among various contexts in order to derive a queue ID.

For example, assume that a stream of data is received at the input buffer 360 from the BQU wherein the stream of data comprises a plurality of contexts, each having a different flow. The hashing function transforms the flow of each context to a queue (Q) ID such that the stream comprises Q1 and Q7 contexts. The input buffer 360 distributes these contexts among the clusters 345 of the processing engine 300 such that Q1 is provided to Cluster 0 and Q7 is provided to Cluster 1. According to conventional FIFO ordering, each context processed by a cluster must complete in the order of distribution by the input buffer to the clusters. That is, if one processor of a cluster stalls, then all processors of all clusters stall in order to maintain the FIFO ordering.

In accordance with an aspect of the present invention, however, assignment of queue IDs to contexts of different flows allows out-of-order processing by the clusters. That is, if the processors of Cluster 1 complete their processing of Q7 before the processors of Cluster 0 finish processing Q1, then Q7 may be forwarded to the output buffer 5 370 and out of (off) the processing engine before Q1. As a result, the order of processing by the engine and transmission off the engine may comprise Q7 and Q1.

The sequence controller further enables selective FIFO ordering particularly in the case when there are more than one contexts associated with a particular flow or queue ID. For example assume that a second context having a flow that transforms to Q1 arrives at 10 the input buffer 360 and is provided to a different cluster (e.g., Cluster 7) for processing. Thus, a first context of Q1 is assigned sequence number “1”, e.g., Q (1,1) and is provided to Cluster 0, while a second context of Q1 is assigned sequence number “2”, e.g., Q (1,2) and is provided to Cluster 7. Assume further that TMC 0 of Cluster 0 stalls during its processing of Q (1,1) or that processing of Q (1,1) consumes a time period that exceeds 15 the duration of a phase. Meanwhile, TMCs 0,1 of Cluster 7 complete their processing of Q (1,2).

According to the selective FIFO ordering aspect of the present invention, context Q (1,2) cannot be transmitted to the output buffer 370 and off the processing engine 300 until processing of context Q (1,1) completes and that context has been transferred from 20 the processing engine. That is, FIFO ordering is enforced for contexts associated with the same flow (or having the same queue ID). The novel control mechanism identifies dependencies among various contexts (i.e., those contexts associated with a particular flow) and imposes ordering among those contexts; however, the sequence control mechanism allows out-of-order processing and completion of those contexts having no dependencies.

25 In the illustrative embodiment, a second circuit of the BQU 210 functions as an incrementor 214 to generate the sequence number associated with a particular context of a flow by, e.g., incrementing a value. According to the invention, sequence numbers are employed to denote an order of contexts within a flow and the BQU assigns those sequence numbers to consecutive contexts in an increasing order. It should be noted, however, that if the input buffer 360 is not configured to interpret a sequence number and/or 30

queue ID assigned by the BQU, it may assign its own sequence number and/or queue ID to contexts of a flow.

Fig. 4 is a schematic block diagram of a context 400 that is constructed by the BQU 210 and provided to the input buffer 360 of the processing engine 300. The context 5 is defined by a predetermined number of bytes of data and includes an appended header 410 having a plurality of fields. Specifically, these fields include a row direction (RD) field 412, a queue ID field 414 and a sequence number field 416. The content of the RD field 412 instructs (i.e., forces) the input buffer 360 to transfer the context to a particular cluster 345 as opposed to the next available cluster of the processing engine. The queue 10 ID field 414 contains a 32-bit queue ID assigned to the context 400 and, if necessary, the sequence number field 416 contains a 16-bit sequence number assigned to the context.

The context 400 is transferred by the BQU to the input buffer 360, which parses (“strips off”) the queue ID and sequence number for storage in a register 362 of the input buffer. The remaining portion of the context (which typically comprises a packet header) 15 is stored in an input buffer portion 336 of the context memory 330 contained within the cluster to which the context is transferred. The information stored in the register 362 is associated with the context 400 stored in the context memory 330; as a result, the register information “flows” with its context through the processing engine.

In the illustrative embodiment, there are preferably four (4) portions of the context memory contained in each cluster of the processing engine: portions 332, 334 for 20 each TMC 0,1, the input buffer portion 336 and an output buffer portion 338. In addition, there is an out-of-band register 335 corresponding to the context memory 330 of each cluster 345 and configured to store the queue ID and the sequence number “control” contents passed from register 362. By separating the control information from the actual context “data” (e.g., packet header), software executing on the TMC processors can easily process the context data without parsing the control information associated with that context. The sequence number and queue ID control information flow with their context 25 and are not actually used again until the output buffer 370 employs that information to determine when the particular context may be transmitted from the processing engine.

According to the invention, the processing engine 300 also includes an input sequence controller 365 coupled to the input buffer 360 and an output sequence controller 375 coupled to the output buffer 370. These controllers, which are preferably implemented as combinational logic circuits, enable out-of-order context processing among the clusters 345 of processing engine 300, while enforcing FIFO synchronization ordering among those clusters for contexts of a flow to thereby reduce undesired processing delays among those processors. The controllers are connected to a data structure 500 that is shared between the controllers. In the illustrative embodiment, this shared data structure is preferably a content addressable memory (CAM) that maintains a list of active flows or “queues” in the processing engine. By utilizing a queue (Q) CAM structure, the sequence controllers can search, in parallel, for all active queues in the processing engine.

Fig. 5 is a schematic block diagram of the QCAM 500 comprising a plurality of entries 510, each containing a 32-bit queue field 512, a 5-bit count field 514 and a 1-bit active field 516. In addition, a 16-bit minimum sequence field 520 may be included within each entry 510 of the QCAM 500 or may be implemented as an entry 520 of another data structure associated with the QCAM. The queue field 512 stores a queue ID of a context associated with a particular queue/flow, whereas the count field 514 stores information indicating the number (i.e., the count) of contexts associated with the flow (or queue ID) that is active in the processing engine and the active field 516 indicates whether the QCAM entry is valid (i.e., whether the queue is active in the processing engine). Notably, a context is “active” in the processing engine if it is stored in a context memory 330 of a cluster 345 between the input and output buffers. The minimum sequence field 520 is used to enforce ordering within a flow having a common queue ID and, to that end, stores the lowest sequence number of a context 400 for a flow/queue that is active in the processing engine. Notably, the value of the minimum sequence number is relative to the number of contexts associated with the flow that is active in the processing engine; this number is incremented by the output sequence controller 375 as contexts are transmitted from the output buffer off the processing engine 300.

In the illustrative embodiment, ordering of contexts associated with a particular flow is enforced at the egress point (e.g., output buffer 370) of the processing engine 300

instead of the ingress point. Enforcement of ordering at the ingress point (e.g., the input buffer 360) of the processing engine requires that the input buffer sequentially issue contexts of the same flow to the same cluster. By using sequence numbers and queue IDs as defined herein, the sequence controller mechanism can issue contexts of the same flow to any cluster, thereby exploiting the parallelism inherent in a systolic array, such as the arrayed processing engine 300. The sequence numbers and queue IDs ensure that transmission of a previous context from the processing engine occurs prior to transmission of a subsequent context associated with that flow.

For example, two contexts having the same queue ID but different sequence numbers, e.g., Q (1,1) and Q (1,2) arrive at the input buffer 360 and are issued to different clusters, e.g., Cluster 0 and 1, respectively, for immediate processing. As described herein, Cluster 1 is unable to transmit context Q (1,2) to the output buffer, even if it completes processing, until Cluster 0 has completed processing of context Q (1,1) and transmitted that context to the output buffer 370. That is, the output buffer ensures that Cluster 1 is not allowed to transmit its context until Cluster 0 has completed processing and transmission of its context to that buffer 370.

In accordance with the present invention, the input and output sequence controllers execute a novel algorithm that ensures that all contexts in the same queue are processed with the same number of phases. The algorithm is apportioned into an input function executable by the input sequence controller 365 and an output function executable by the output sequence controller 375. Illustrative program code segments representative of the novel input and output functions are as follows:

Input Function:

```

if (Q[inputContext.queueID].isActive) {
    25   Q[inputContext.queueID].count++;
} else {
    Q[inputContext.queueID].isActive = TRUE;
    Q[inputContext.queueID].count = 1;
    Q[inputContext.queueID].minSeq = inputContext.sequence; //could be set to 1
    30
}

```

Output Function:

```

if (Q[outputContext.queueID].minSeq == outputContext.sequence) {
    TransmitRequest for row N;
    {{ // operation performed in 1 cycle
5      Q[inputContext.queueID].minSeq++;
      Q[inputContext.queueID].count--;
      if (Q[inputContext.queueID].count == 0)
          Q[inputContext.queueID].isActive = FALSE;
    }}
10 } else {
      STALL
}

```

The notation $Q[xContext.queueID]$ denotes a lookup operation into the QCAM on the basis of the queue ID. Referring to the input function code, if a lookup operation into the QCAM 500 by the input sequence controller 365 results in a “hit” on an entry 510 having the queue ID of a particular context, then the active and count fields of that entry are examined. Specifically, if a context 400 having a particular queue ID is active in the processing engine (i.e., the active bit is asserted) and a subsequent “input” context of the same queue ID arrives at the input buffer 360, then the input function specifies incrementing the count. In other words, the count information stored in the count field 514 of the QCAM is incremented for each successive context of the queue ID.

If, however, the lookup operation into the QCAM results in a “miss”, a new entry 510 of the QCAM is allocated and the queue ID associated with the flow of the input context is inserted into the queue field 512 of that entry. Thereafter, the active field 516 of the allocated QCAM entry is activated (e.g., set to 1) and the count field 514 is incremented (e.g., set to 1) indicating that the input context is the first active context in the processing engine for the queue. In addition, the content of the minimum sequence number field 520 associated with the allocated queue entry 510 is set to the assigned sequence number of the first context of a flow. For example, a first context that arrives at the input buffer having a particular flow is assigned a queue ID for that flow and a sequence num-

ber equal to, e.g., “1”. The minimum sequence number is thus set to “1”, even though the next context arriving at the input buffer having the same queue ID is assigned a sequence number of “2”.

When the BQU 210 transmits a context 400 to the input buffer, the input buffer
 5 360 issues a “new context” control signal over line 364 to the input sequence controller 365 that triggers execution of the input function by that controller. Essentially, the new context signal informs the input sequence controller that a new input context is available within the processing engine and also instructs the controller 365 to update the QCAM 500 with appropriate information. The appropriate information preferably includes the
 10 sequence number and queue ID associated with that new context. The input buffer 360 transmits the sequence number and queue ID signals over lines 366, 368, respectively, to the input sequence controller.

It should be noted that the input sequence controller 365 executes the input function once for each input context 400 received at the input buffer 360. On the other hand,
 15 the output sequence controller 375 executes the output function for each row/cluster 345 that requests transmission of its processed “output” context to the output buffer 370. Thus, for the illustrative processing engine 300 having eight (8) clusters 345, the output sequence controller 375 executes the output function 8 times, one for each cluster.

Upon completing processing of a context, a processor, e.g., TMC 1, of a cluster
 20 345 issues a *pdone* completion signal over line 352 to the output buffer 370. In response, the output buffer issues a context request signal over line 372 to the output sequence controller 375 informing the controller that this particular output context is a candidate for transmission off the processing engine. The output buffer also provides to the controller 375 a sequence number signal associated with the output context over line 374 and a queue ID signal associated with that output context over line 376. The output sequence controller grants the request for transmission off the engine by issuing an acknowledge (Ack) signal over line 378 for each context request issued by the output buffer. Each of the signals exchanged between the output buffer and the output sequence controller is replicated, e.g., 8 times, once for each cluster.

Before the output buffer may transmit the output context off the processing engine, the output sequence controller executes the output function to validate that the sequence of that context is correct. That is, the output function is executed by the output sequence controller 365 for each cluster to validate the orderly sequence of output contexts for a particular flow. At this time, the context 400 may be processed out-of-order depending upon its dependencies and associations with other contexts of the particular flow.

Refer now to the output function code. In response to the context request, the output sequence controller 375 scans the QCAM 500 for an entry 510 having the queue ID associated with that context request. When a matching entry is found, the output sequence controller 375 determines whether to grant the transmission request by comparing the sequence number associated with the output context with the minimum sequence number stored in the QCAM for the entry having a matching queue ID. If these sequence numbers are equal, the output context may be transmitted from the cluster (row N) to the output buffer. Upon rendering the decision to transmit the context from the processing engine, the sequence controller 374 performs the following update operations on the corresponding entry 510 of the QCAM 500. Specifically, the output sequence controller 375 (i) increments the minimum sequence number associated with the queue ID of the output context, (ii) decrements the count associated with that context, (iii) tests the count to be zero and, if the count equals zero, (iv) deasserts the active bit associated with the entry to denote that the queue is no longer active in the processing engine.

On the other hand, if the minimum sequence number does not equal the output context sequence number, then an output context that is out-of-order within a particular flow has completed processing and has requested transmission from the processing engine. In other words, the output context is “out-of-order” with respect to at least one other active context in the processing engine and, as a result, the transmission request is not granted for the output context. For example, assume there are two contexts having the same queue ID that are active within the processing engine. The first context issued from that queue takes a relatively long time to complete processing at, e.g., Cluster 0. Even though a second context from that queue (flow) is issued to another cluster, e.g.,

Cluster 7, at a later time, it completes processing sooner than the first context. According to the invention, the request for transmission of the second context from the processing engine is not granted until the first context associated with that flow is processed and transmitted off the engine.

5 Referring to the output function code segment, the example above manifests as a failed condition of whether the minimum sequence number equals the output context sequence number and, therefore, the output function transitions to a STALL state for that particular output context sequence number. For example, if the minimum sequence number is “1” and the output context sequence number of the second context is “2”, the condition fails because “1” does not equal “2”. Essentially, the output buffer 370 and output sequence controller 375 must wait for the first context to complete and, as a result, 10 transmission of the second context off the processing engine is stalled.

In accordance with the present invention, execution of the output function ensures and enforces order at an output interface of the processing engine for those contexts associated with a particular flow. That is, the second context associated with the flow continues to stall until the first context associated with that flow completes processing. Once 15 the first context completes processing, the output buffer requests transmission of that processed context from the processing engine and the output sequence controller executes the output function for the processed context. Execution of the output function indicates that the first context of the flow meets the condition for transmission and, thus, allows 20 that context to be transmitted from the processing engine. Once the update operations of the output function are executed by the output sequence controller, the second context having the particular queue ID is allowed to transmit off the processing engine. This is an example of how order is maintained for contexts associated with a particular flow in 25 accordance with the selective FIFO ordering technique of the present invention.

In the illustrative embodiment, the output buffer 360 continues to assert the context request 372 for transmission of a context until the request is acknowledged by the output sequence controller 375. That is, until an Ack signal 378 is provided to acknowledge the transmission request for a context having a particular sequence number and 30 queue ID, that context remains in a stalled state of the output function. Every transmit

request issued by the output buffer to the output sequence controller must be acknowledged prior to allowing the context to transmit off the processing engine.

Operation of the processing engine in accordance with the improved control mechanism will now be described with respect to Figs. 1-5. When a context 400 is received at the input buffer, the buffer 360 parses the queue ID and sequence number, and loads that information into register 362. The input buffer provides the sequence number and queue ID to the input sequence controller 365 so that it can update an appropriate entry 510 of the QCAM 500. The input buffer 360 then loads the context 400 into the context memory 330 of a particular cluster 345 and the context proceeds to “ping-pong” through the portions of the memory 330 as it is processed by the processors 350. Meanwhile, the queue ID and sequence number are passed to register 335 associated with the context memory of the particular cluster.

The context 400 then propagates to the output buffer portion 338 of the context memory. The output buffer determines that it has an active context identified by the sequence number and queue ID information stored in register 335 that has completed processing and that requests transmission off the engine. Accordingly, the output buffer 370 transfers that information to the output sequence controller 375 in accordance with a context request for transmission of the output context from the processing engine. In response, the output sequence controller 375 executes the output function to determine whether that particular context may be transmitted off of the engine. Based upon the states of the other clusters, the controller 375 acknowledges the context request and sends that acknowledgement to the output buffer 370. Thus, the context is granted access to the output interface of the processing engine.

The contents of the register 335, along with context 400, are transmitted from the output buffer over the output interface from the processing engine 300 to the BQU 210. If the context is not allowed to transmit off the processing engine because it is out of sequence, transmission is stalled until the dependencies are resolved. Thus, the incoming order for a particular queue is maintained as the outgoing order for that particular queue in accordance with a “queue ordering” mode, as described herein. In an alternate embodiment, the BQU may maintain a sequential order of every context issued to the proc-

essing engine by assigning a fixed queue ID and incrementing the sequence number for each issued context.

While there has been shown and described an illustrative embodiment of a sequence controller configured to implement out-of-order processing of contexts by processors of an arrayed processing engine, it is to be understood that various other adaptations and modifications may be made within the spirit and scope of the invention. For example, the selective FIFO ordering technique may provide an advantage with even two clusters/rows of the processing engine. That is, assume that a context of a particular flow takes ten (10) times longer to process than other contexts of other flows. The context requiring substantially long processing is loaded into a first row of a processing engine having two rows of processors. Nine (9) other contexts of other flows may be processed by the second row of the processing engine during the time that the first context is processed by the first row. Thus, selective FIFO ordering facilitates out-of-order processing for an engine having a plurality of clusters/rows wherein there are no dependencies upon the contexts processed by the processors of those rows.

The foregoing description has been directed to specific embodiments of this invention. It will be apparent, however, that other variations and modifications may be made to the described embodiments, with the attainment of some or all of their advantages. Therefore, it is the object of the appended claims to cover all such variations and modifications as come within the true spirit and scope of the invention.

What is claimed is:

CLAIMS

1. A method for enabling out-of-order processing of contexts by processors of a multi-
2 processor system, the processors arrayed as a plurality of clusters embedded between in-
3 put and output buffers, the method comprising the steps of:

4 assigning each context a queue identifier (ID) and a sequence number, the queue
5 ID uniquely identifying a flow of the context and the sequence number denoting an order
6 of the context within the flow;

7 distributing the contexts from the input buffer to the clusters;

8 allowing out-of-order context processing among the clusters for contexts having
9 different queue IDs; and

10 enforcing first in, first out (FIFO) synchronization context processing among the
11 clusters for contexts having the same queue ID.

1. 2. The method of Claim 1 wherein the step of assigning comprises the step of deriving
2 the queue ID using information that enables identification of dependencies among the
3 contexts.

1. 3. The method of Claim 1 wherein the step of generating comprises the step of trans-
2 forming flow parameters of a context to the queue ID in accordance with a hash function.

1. 4. The method of Claim 1 wherein the step of assigning comprises the step of increment-
2 ing a predetermined value to generate the sequence number.

1. 5. The method of Claim 1 further comprising the step of coupling an input sequence con-
2 troller to the input buffer and an output sequence controller to the output buffer.

1. 6. The method of Claim 5 further comprising interconnecting the input and output se-
2 quence controllers with a data structure that maintains a list of active flows in the system.

- 1 7. The method of Claim 6 wherein the data structure is a content addressable memory
2 (CAM) having a plurality of entries.
- 1 8. The method of Claim 6 further comprising the step of providing a queue field and a
2 minimum sequence field within each entry of the data structure.
- 1 9. The method of Claim 8 further comprising the step of executing an input function at
2 the input sequence controller to update the data structure with the sequence number and
3 queue ID associated with a new context.
- 1 10. The method of Claim 9 wherein the step of updating comprises the steps of:
2 storing the queue ID in the queue ID field of an appropriate entry; and
3 storing a lowest sequence number of a context for a flow that is active in the sys-
4 tem in the minimum sequence field of the entry.
- 1 11. The method of Claim 10 wherein the step of storing a lowest sequence number com-
2 prises the step of setting the content of the minimum sequence field to the assigned se-
3 quence number of the first context of a flow.
- 1 12. The method of Claim 8 further comprising the step of executing an output function at
2 the output sequence controller to validate one of the out-of-order processing and FIFO
3 synchronization processing of the contexts.
- 1 13. Apparatus for enabling out-of-order processing of contexts by processors of a proc-
2 essing engine, the processors arrayed as a plurality of clusters, the apparatus comprising:
3 a hash function adapted to transform flow parameters of a context to a queue iden-
4 tifier (ID) that uniquely identifies a flow of the context;

5 an incrementor coupled to the hash function and configured to increment a prede-
6 termined value to generate a sequence number denoting an order of the context within the
7 flow;

8 an input buffer of the processing engine coupled to the hash function and incre-
9 mentor, the input buffer distributing the contexts to the clusters; and

10 a sequence control mechanism that allows out-of-order context processing among
11 the clusters for contexts having different queue IDs and enforces first in, first out (FIFO)
12 synchronization context processing among the clusters for contexts having the same
13 queue ID.

1 14. The apparatus of Claim 13 wherein the sequence control mechanism comprises an
2 input sequence controller coupled to the input buffer and an output sequence controller
3 coupled to an output buffer of processing engine.

1 15. The apparatus of Claim 14 wherein the sequence control mechanism further com-
2 prises a data structure coupled between the input and output sequence controllers, the
3 data structure maintaining a list of active flows in the system.

1 16. The apparatus of Claim 15 wherein the data structure is a content addressable mem-
2 ory (CAM) having a plurality of entries, each entry including a queue field that stores the
3 queue ID of a context and a minimum sequence field that stores a lowest sequence num-
4 ber of a context for a flow that is active in the engine.

1 17. A computer readable medium containing executable program instructions for ena-
2 bling out-of-order processing of contexts by processors of a processing engine, the proc-
3 essors arrayed as a plurality of clusters embedded between input and output buffers, the
4 executable program instructions comprising program instructions for:

5 assigning each context a queue identifier (ID) and a sequence number, the queue
6 ID uniquely identifying a flow of the context and the sequence number denoting an order
7 of the context within the flow;

8 distributing the contexts from the input buffer to the clusters;

9 allowing out-of-order context processing among the clusters for contexts having
10 different queue IDs; and

11 enforcing first in, first out (FIFO) synchronization context processing among the
12 clusters for contexts having the same queue ID.

1 18. The computer readable medium of Claim 17 further comprising program instructions
2 for executing an input function at an input sequence controller coupled to the input
3 buffer, the input function updating a data structure with the sequence number and queue
4 ID associated with a new context, the data structure maintaining a list of active flows in
5 the processing engine.

1 19. The computer readable medium of Claim 18 further comprising program instructions
2 for executing an output function at an output sequence controller coupled to the output
3 buffer, the output function validating one of the out-of-order processing and FIFO syn-
4 chronization processing of the contexts.

1 20. A method for enabling out-of-order processing of contexts by processors of a multi-
2 processor system, the processors arrayed as a plurality of clusters embedded between in-
3 put and output buffers, the method comprising the steps of:

4 assigning each context a queue identifier (ID) and a sequence number, the queue
5 ID uniquely identifying a flow of the context and the sequence number denoting an order
6 of the context within the flow;

7 providing the queue ID and sequence number to an input sequence controller cou-
8 pled to the input buffer;

9 updating a data structure with the queue ID and sequence number at the input se-
10 quence controller, the data structure maintaining a list of active flows in the system;

11 processing the context at the processors of the cluster;

12 at an output sequence controller coupled to the output buffer allowing one of out-
13 of-order processing and first in, first out synchronization processing of the context de-
14 pending upon the queue ID and sequence number of the context.

ABSTRACT OF THE DISCLOSURE

A sequence control mechanism enables out-of-order processing of contexts by processors of a symmetric multiprocessor system having a plurality of processors arrayed as a processing engine. The processors of the engine are preferably arrayed as a plurality of rows or clusters embedded between input and output buffers, wherein each cluster of processors is configured to process contexts in a first in, first out (FIFO) synchronization order. However, the sequence control mechanism allows out-of-order context processing among the clusters of processors, while selectively enforcing FIFO synchronization ordering among those clusters on an as needed basis, i.e., for certain contexts. As a result, the control mechanism reduces undesired processing delays among those processors.

1/4




FIG. 1

2/4




FIG. 2




FIG. 3

4/4




FIG. 4




FIG. 5

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below-named inventor, I hereby declare that:

My residence, post-office address, and citizenship are as stated below next to my name.

I believe I am an original, first, and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled SEQUENCE CONTROL MECHANISM FOR ENABLING OUT OF ORDER CONTEXT PROCESSING, the specification of which is attached hereto and identified by Cesari and McKenna File No. 112025-0196.

I hereby state that I have reviewed and understand the contents of the above-identified application specification, including the claims, as amended by any amendment specifically referred to herein.

I acknowledge the duty to disclose all information known to me that is material to patentability in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code §119(a)-(d) of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate filed by me on the same subject matter having a filing date before that of the application on which priority is claimed: None.

I hereby claim the benefit under Title 35, United States Code §119(e) of the following U.S. provisional application: None.

I hereby claim the benefit under Title 35, United States Code §120, of the United States Application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United State Code, §112, I acknowledge the duty to disclose all information that is material to patentability in accordance with Title 37, Code of Federal Regulations, §1.56, and which became available to me between the filing date of the prior application and the national or PCT international filing date of this application: None.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment or both under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint Michael E. Attaya, Reg. No. 31,731; Charles J. Barbas, Reg. No. 32,959; Joseph H. Born, Reg. No. 28,283; John L. Capone, Reg. No. 41,656; Robert A. Cesari, Reg. No. 18,381; Yong S. Choi, Reg. No. 43,324; Brian C. Dauphin, Reg. No. 40,983; Steven J. Frank, Reg. No. 33,497; Christopher K. Gagne, Reg. No. 36,142; A. Sidney Johnston, Reg. No. 29,548; William A. Loginov, Reg. No. 34,863; John F. McKenna, Reg. No. 20,912; Rama B. Nath, Reg. No. 27,072; Martin J. O'Donnell, Reg. No. 24,204; Thomas C. O'Konski, Reg. No. 26,320; Edwin H. Paul, Reg. No. 31,405; Michael R. Reinemann, Reg. No. 38,280; Rita M. Rooney, Reg. No. 30,585; Heather B. Shapiro, Reg. No. 41,305; Patricia A. Sheehan, Reg. No. 32,301; and Joseph Stecewycz, Reg. No. 34,442, Cesari and McKenna, LLP, 30 Rowes Wharf, Boston, Mass. 02110, jointly, and each of them severally, my attorneys and attorney, with full power of substitution, delegation and revocation, to prosecute this application, to make alterations and amendments therein, to receive the patent and to transact all business in the Patent and Trademark Office connected therewith. Please direct all telephone calls to Charles J. Barbas at (617) 951-2500. Please address all correspondence to Charles J. Barbas.

Darren Kerr	Date	
Residence:	853 Forest Ave. Palo Alto, CA 94301	
Citizenship	United States of America	
Post Office Address:	Same as above	
<u>Jeffery B. Scott</u>		7/24/00
Residence:	901 Limerick Lane Durham, NC 27713	

Citizenship	United States of America
Post Office Address:	Same as above

John William Marshall 7/24/00
John William Marshall Date

Residence: 212 Preston Pines Drive
Cary, North Carolina 27513

Citizenship United States of America

Post Office Address: Same as above

Kenneth H. Potter 7/24/00
Kenneth H. Potter Date

Residence: 5404 Amsterdam Place
Wake County
Raleigh, NC 27606

Citizenship U.S.A.

Post Office Address: Same as above

Scott Nellenbach 7/24/00
Scott Nellenbach Date

Residence: 208 Burgwin Wright Way
Apex, North Carolina 27502

Citizenship United States of America

Post Office Address: Same as above

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below-named inventor, I hereby declare that:

My residence, post-office address, and citizenship are as stated below next to my name.

I believe I am an original, first, and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled SEQUENCE CONTROL MECHANISM FOR ENABLING OUT OF ORDER CONTEXT PROCESSING, the specification of which is attached hereto and identified by Cesari and McKenna File No. 112025-0196.

I hereby state that I have reviewed and understand the contents of the above-identified application specification, including the claims, as amended by any amendment specifically referred to herein.

I acknowledge the duty to disclose all information known to me that is material to patentability in accordance with Title 37, Code of Federal Regulations, §1.56.


I hereby claim foreign priority benefits under Title 35, United States Code §119(a)-(d) of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate filed by me on the same subject matter having a filing date before that of the application on which priority is claimed: None.

I hereby claim the benefit under Title 35, United States Code §119(e) of the following U.S. provisional application: None.

I hereby claim the benefit under Title 35, United States Code §120, of the United States Application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United State Code, §112, I acknowledge the duty to disclose all information that is material to patentability in accordance with Title 37, Code of Federal Regulations, §1.56, and which became available to me between the filing date of the prior application and the national or PCT international filing date of this application: None.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment or both under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint Michael E. Attaya, Reg. No. 31,731; Charles J. Barbas, Reg. No. 32,959; Joseph H. Born, Reg. No. 28,283; John L. Capone, Reg. No. 41,656; Robert A. Cesari, Reg. No. 18,381; Yong S. Choi, Reg. No. 43,324; Brian C. Dauphin, Reg. No. 40,983; Steven J. Frank, Reg. No. 33,497; Christopher K. Gagne, Reg. No. 36,142; A. Sidney Johnston, Reg. No. 29,548; William A. Loginov, Reg. No. 34,863; John F. McKenna, Reg. No. 20,912; Rama B. Nath, Reg. No. 27,072; Martin J. O'Donnell, Reg. No. 24,204; Thomas C. O'Konski, Reg. No. 26,320; Edwin H. Paul, Reg. No. 31,405; Michael R. Reinemann, Reg. No. 38,280; Rita M. Rooney, Reg. No. 30,585; Heather B. Shapiro, Reg. No. 41,305; Patricia A. Sheehan, Reg. No. 32,301; and Joseph Stecewycz, Reg. No. 34,442, Cesari and McKenna, LLP, 30 Rowes Wharf, Boston, Mass. 02110, jointly, and each of them severally, my attorneys and attorney, with full power of substitution, delegation and revocation, to prosecute this application, to make alterations and amendments therein, to receive the patent and to transact all business in the Patent and Trademark Office connected therewith. Please direct all telephone calls to Charles J. Barbas at (617) 951-2500. Please address all correspondence to Charles J. Barbas.



Darren Kerr

6/13/00

Date

Residence: 853 Forest Ave.
Palo Alto, CA 94301

Citizenship United States of America

Post Office Address: Same as above

Jeffery B. Scott

Date

Residence: 901 Limerick Lane
Durham, NC 27713

Citizenship United States of America

Post Office Address: Same as above

John William Marshall	Date
Residence:	212 Preston Pines Drive Cary, North Carolina 27513
Citizenship	United States of America
Post Office Address:	Same as above

Kenneth H. Potter	Date
Residence:	5404 Amsterdam Place Wake County Raleigh, NC 27606
Citizenship	U.S.A.
Post Office Address:	Same as above

Scott Nellenbach	Date
Residence:	208 Burgwin Wright Way Apex, North Carolina 27502
Citizenship	United States of America
Post Office Address:	Same as above